arXiv:1704.03113v1 [physics.comp-ph] 11 Apr 2017

Molecular geometry and vibrational frequencies by parallel sampling

Jan Vrbik[1]

[1]*Department of Mathematics and Statistics, Brock University, Canada*

April 12, 2017

Abstract

Quantum Monte Carlo is an efficient technique for finding the ground-state energy and related properties of small molecules. A major challenge remains in accurate determination of a molecule's geometry, i.e. the optimal location of its individual nuclei and the frequencies of their vibration. The aim of this article is to describe a simple technique to accurately establish such properties. This is achieved by varying the trial function to accommodate changing geometry, thereby removing a source of rather unpleasant singularities which arise when the trial function is fixed (the traditional approach).

Keywords: QMC; molecular geometry; vibrational frequencies; energy derivatives

1 Introduction

In this article we describe a new technique of PARALLEL SAMPLING for computing various properties (those related to location and vibration of nuclei in particular) of small molecules. But first we present (for the benefit of a non-specialist) a brief review of general Quantom Monte Carlo.

To find the GROUND-STATE ENERGY of a molecule given its GEOMETRY (i.e. distances between its individual nuclei, for this purpose considered fixed – this constitutes the so-called CLAMPED approximation), one has to solve the following eigenvalue problem

$$-\frac{1}{2}\nabla^2\Phi(\mathbf{x})+V(\mathbf{x})\Phi(\mathbf{x})=\mathcal{E}\,\Phi(\mathbf{x}) \tag{1}$$

where $\Phi(\mathbf{x})$ is an (unknown) function of the locations of the molecule's (say n) *electrons* collectively denoted $\mathbf{x}$, ∇^2 is the corresponding$3n$-dimensional Laplace operator, $V(\mathbf{x})$ is the

molecule's electro-static potential and $\mathcal{E}$ is the corresponding *lowest* eigenvalue. The trouble is that this equation can be solved only approximately (usually using the VARIATIONAL PRINCIPLE) by a function we denote $\Psi(\mathbf{x})$ and call the TRIAL FUNCTION. In this approximate solution, $\mathcal{E}$ itself becomes a function of $\mathbf{x}$ (the so-called local energy, defined shortly) which is only approximately 'constant'. We can improve the accuracy of such a solution by increasing the complexity of $\Psi(\mathbf{x})$, correspondingly increasing the challenge of optimizing its parameters.

Having such variational solution (of whatever quality), Diffusion Monte Carlo aims to substantially improve the corresponding estimates of $\mathcal{E}$ (and eventually of its derivatives) by the following procedure (see, for example [1] and [10]):

1. Based on $\Psi(\mathbf{x})$, it defines two more functions, namely

   (a) the so-called DRIFT function

   $$\mathbf{F}(\mathbf{x}) := \frac{\nabla\ \Psi(\mathbf{x})}{\Psi(\mathbf{x})} \tag{2}$$

   where ∇ represents the vector of all partial derivatives (with respect to each component of $\mathbf{x}$) – note that this makes $\mathbf{F}(\mathbf{x})$ a *vector* with $3n$ components

   (b) and the LOCAL ENERGY

   $$E(\mathbf{x}) := -\frac{1}{2}\frac{\nabla^2\Psi(\mathbf{x})}{\Psi(\mathbf{x})} + V(\mathbf{x}) \tag{3}$$

   which is a scalar function of$\mathbf{x}$. We mention in passing that

   $$\frac{E(\mathbf{x})\Psi(\mathbf{x})^2 d\mathbf{x}}{\Psi(\mathbf{x})^2 d\mathbf{x}} \tag{4}$$

   represents the so-called VARIATIONAL approximation to $\mathcal{E}$ (always higher than the correct value), where the integration is over all space.

2. Starting with a rather arbitrarily chosen collection (called an ENSEMBLE, which should be as large as our computer can handle) of CONFIGURATIONS (numerical values of $\mathbf{x}$), and using a small value of τ (called TIME STEP), we then advance each configuration $\mathbf{x}$ to a new $\mathbf{x}'$ by

   $$\mathbf{x}' = \mathbf{x} + \tau\ \mathbf{F}(\mathbf{x}) + \chi \tag{5}$$

   where χ is a *random* vector with$3n$ components, drawn independently from the Normal distribution with the mean of 0 and the standard deviation of $\sqrt{\tau}$. The $\tau\ \mathbf{F}(\mathbf{x})$ part of this 'move' is the actual 'drift', while the χ part is referred to as DIFFUSION. We then compute, for each resulting new configuration, the value of $E(\mathbf{x}')$; this constitutes a so-called ITERATION.

3. We then repeat as many of these iterations as possible ($\mathbf{x}'$ always becomes the 'old' $\mathbf{x}$ of the next iteration). For each configuration, we also keep track of the sum of all (current and past) values of the local energy, *depreciated* by the factor of

$$1 - \frac{a\ \tau}{1 - \ln \tau} \tag{6}$$

where a is an empirically chosen constant. To be more specific: in each iteration the old sum is updated by multiplying it by (6) and adding to it the new value of $E(\mathbf{x}')$; the result will be called $S(\mathbf{x}')$. We then compute the iteration's *weighted* average of the current $E(\mathbf{x}')$ values by

$$\overline{E} := \frac{\sum E(\mathbf{x}') \exp(-\tau\ S(\mathbf{x}'))}{\sum \exp(-\tau\ S(\mathbf{x}'))} \tag{7}$$

where the summation is over all configurations of the ensemble and

$$W(\mathbf{x}') := \exp(-\tau\ S(\mathbf{x}')) \tag{8}$$

are the corresponding weights. Note that the simple average of the $E(\mathbf{x}')$ values would estimate (4) instead; this is true in general: whenever we replace $W(\mathbf{x}')$ by 1 (which corresponds to *regular* averaging) we get the so-called (and rather inferior) *variational* estimates – this applies to our subsequent estimation of various derivatives of $\mathcal{E}$ (see [12]) as well.

4. By monitoring the values of $\overline{E}$ from iteration to iteration, we can easily notice that after a while the process stabilizes (the values of $\overline{E}$ no longer systematically increase or decrease; they only randomly fluctuate around some mean value). At the end of the simulation (after thousands of iterations) we again average (this time in the regular manner) the iteration averages themselves (but only the EQUILIBRATED ones, after stabilization), denoting the resulting GRAND MEAN by $\overline{\overline{E}}(\tau)$.

5. Even this $\overline{\overline{E}}(\tau)$ is not yet our final estimate of $\mathcal{E}$, since it contains a τ-related BIAS. To remove it, we have to repeat the simulation with several *different* values of τ (three to five such values are usually sufficient) and extrapolate the corresponding grand means to $\tau = 0$ by a simple (usually quadratic) regression.

We should mention (since it is far from obvious) that this corresponds to replacing, in (4), $\Psi(\mathbf{x})^2$ by the so-called MIXED distribution, proportional to the *product* of the exact solution $\Phi(\mathbf{x})$ and the variational solution $\Psi(\mathbf{x})$. Since $\Phi(\mathbf{x})$ is thus forced to have the same nodes as $\Psi(\mathbf{x})$, this introduces a small (in our case almost negligible) FIXED-NODE error into the final estimate of $\mathcal{E}$. Subsequently (and due to this error somehow incorrectly), we call the corresponding results *exact*, those using regular averaging in (7) are called *variational*.

Note that the proposed technique has no need for the so-called PURE sampling from a distribution proportional to $\Phi(\mathbf{x})^2$; such sampling does *not* properly estimate $\mathcal{E}$ and is useful only in the context of Hellmann-Feynman Theorem.

2 Parallel simulation

We now consider the case of $V(\mathbf{x})$, and consequently of $\Phi(\mathbf{x})$ and $\mathcal{E}$, being functions of a parameter λ (the technique can be easily extended to deal with *several* such parameters). Usually λ represents the strength of an external perturbation (such as a uniform electrical field), but it can also be the distance between two of the molecule's nuclei (as mentioned already, these distances constitute what we call the molecule's geometry). We assume that the trial solution $\Psi(\mathbf{x})$ is *also* a function of λ (e.g. atomic orbitals of the trial function are centered on individual nuclei, and move with these). This implies that both $\mathbf{F}(\mathbf{x})$ and $E(\mathbf{x})$ also vary with λ (now, an implicit argument of both functions).

The objective of this article is to find a good approximation to the first, second and perhaps higher *derivatives* of the ground-state energy with respect to λ (at $\lambda = 0$ in the case of external perturbations, or at $\lambda = \lambda_0$, where λ_0 represents the initial inter-nuclear distance). We want to emphasize that such derivatives yield *most* of the basic molecular properties, e.g. the dipole moment, polarizability, harmonic and unharmonic frequencies etc., including the molecule's *optimal* geometry.

We achieve this goal by a modification of the so-called CORRELATED SAMPLING of [11] and [8]; instead of using small but finite separation between different values of λ (as they have done), we propose to make these differences truly 'infinitesimal' (which is the main thrust of this article). This alleviates the problem of sudden divergence of parallel configurations (as discussed in [2] and [3]), and eliminates the need not only for subsequent finite-step differentiation but, most importantly, for the so-called space-wrapping of electrons (see [11]). We also deliberately avoid employing the Hellmann-Feynman theorem (recommended only for high-quality trial functions, e.g. those of [5] and [7]), since its application may seriously bias our results (see [6]).

We now describe a procedure to accurately estimate the ground-state energy and its *first* and *second* derivatives with respect to λ; extending this to the case of higher derivatives is then fairly routine.

To accomplish the former, one has to perform the five steps of the previous section, with the following extensions:

1. Throughout the simulation, each configuration $\mathbf{x}$ splits into three (*initially* identical 'companions'), which are then advanced 'in parallel' using thee *slightly* different values of λ (namely $\lambda_0 - \varepsilon$, λ_0, and $\lambda_0 + \varepsilon$) in the drift part of (5), but the *same* random vector χ in the diffusion part. The value of ε (aiming for *practically* 'infinitesimal') is chosen to decrease with τ in the following manner

$$\varepsilon = b\sqrt{\tau} \tag{9}$$

where b is a suitable constant (to keepε in the 10^{-5} to 10^{-4} range).

2. Similarly, the$E(\mathbf{x})$ and$S(\mathbf{x})$ are each computed in the corresponding triplicate (referring to them as E_-, E_0, and E_+, and S_-, S_0 and S_+). The iteration estimates of the ground-state energy $\mathcal{E}$ and its first and second λ derivatives (at $\lambda = \lambda_0$) are computed by

$$\overline{E}_0 \tag{10}$$

$$\frac{\overline{E}_+ - \overline{E}_-}{2\varepsilon} \tag{11}$$

$$\frac{\overline{E}_+ - 2\overline{E}_0 + \overline{E}_-}{\varepsilon^2} \tag{12}$$

respectively, where the averaging (implied by the bar) is done in the manner of (7). Note that each of the three averages uses its own set of weights (slightly different from the other two), namely $\exp(-\tau S_-)$, $\exp(-\tau S_0)$, and $\exp(-\tau S_+)$; from now on, we denote these W_-, W_0, and W_+ respectively.

3. After thousands of iterations, we convert each of these into the corresponding *grand-mean* estimate (by ordinary averaging). Repeating with several different values of τ and extrapolating to $\tau = 0$ yields the ultimate estimates (and their standard errors) of the three quantities.

An important issue to resolve is whether the three nearly parallel paths of any such 'triplet' of configurations will remain at ε – proportional distances throughout the simulation, or whether their separation can slowly but steadily increase. Luckily (and perhaps somehow surprisingly), it appears that this process is rather stable and the probability of any of the three companions diverging from the other two is extremely small. But it *can* happen, and when it does (even to a single triplet of an ensemble of thousands), the damage is large and permanent (the tight bond between originally parallel configurations will never be restored – they now go their own separate ways).

To minimize the probability of this happening, and to expediently fix it when it does happen, we propose to reduce the distance from the 'leading' configuration (the one with $\lambda = \lambda_0$, say $\mathbf{x}_0$) to each of its two companions (say $\mathbf{x}_+$ and $\mathbf{x}_-$) by a factor similar to (6); that means that, at the end of every iteration (and for each triplet of the ensemble). $\mathbf{x}_+$ is replaced by

$$\mathbf{x}_0 + (\mathbf{x}_+ - \mathbf{x}_0) \cdot \left(1 - \frac{a'\,\tau}{1 - \ln \tau}\right) \tag{13}$$

and similarly for $\mathbf{x}_-$.This not only alleviates the problem, but also reduces (quite significantly) the statistical error of all the estimates, at the cost of introducing yet anotherτ-proportional bias (which is then automatically removed by extrapolating to $\tau = 0$). Since this also effectively reduces the value ofε in (11) and (12) by the factor of

$$r := 1 - \frac{a'\,\tau}{1 - \ln \tau} \tag{14}$$

one should adjust the two formulas accordingly; it helps reducing the corresponding τ-related bias.

The a' constant needs to be adjusted empirically; making it too small leaves us with substantial statistical error, but make it too large and the τ -related bias thus introduced becomes impossible to remove by extrapolation.

Note: To find corresponding estimates of the 3rd and 4th derivative, we have to add two extra parallel companions, one with $\lambda = \lambda_0 - 2\varepsilon$ and the other with $\lambda = \lambda_0 + 2\varepsilon$ (let us call the corresponding iteration means of the local energy $\overline{E}_{--}$ and $\overline{E}_{++}$ respectively). The iteration estimates of the two derivatives are then (incorporating the r factor)

$$\frac{\overline{E}_{++} - 2\overline{E}_{+} + 2\overline{E}_{-} - \overline{E}_{-}}{2\varepsilon^3 r^3} \tag{15}$$

and

$$\frac{\overline{E}_{++} - 4\overline{E}_{+} + 6\overline{E}_{0} - 4\overline{E}_{-} + \overline{E}_{--}}{\varepsilon^4 r^4} \tag{16}$$

respectively; the rest of the procedure remains the same (only ε must then be kept, more conservatively, in the 10^{-3} to 10^{-2} range).

2.1 Example (harmonic oscillator)

In our first 'toy' example (taken from [12]) we consider a one-dimensional quantum harmonic oscillator of a unit mass and the spring constant given by a (rather arbitrarily chosen) function of λ, namely

$$\frac{1+\lambda}{1-\lambda} \tag{17}$$

Analytically solving the corresponding eigenvalue equation

$$-\frac{1}{2}\frac{\mathrm{d}^2}{\mathrm{d}x^2}\Phi(x) + \frac{x^2}{2}\frac{1+\lambda}{1-\lambda}\Phi(x) = \mathcal{E}\,\Phi(x) \tag{18}$$

for the lowest eigenvalue yields

$$\mathcal{E} = \frac{1}{2}\sqrt{\frac{1+\lambda}{1-\lambda}} \tag{19}$$

where

$$\Phi(x) \propto \exp(-\mathcal{E}\,x^2) \tag{20}$$

implying that, at $\lambda = 0$, the ground-state energy and its first four λ-derivatives are equal to 0.5, 0.5, 0.5, 1.5, and 4.5 respectively.

Using the trial function

$$\Psi(x) = \left(1 + \frac{2}{5}\frac{1+\lambda}{1-\lambda}x^4\right)^{-1} \tag{21}$$

the variational (and thus visibly inaccurate) solution for the ground-state energy becomes $\sqrt{10}/3 \simeq 1.0541$ times larger than its exact value (19), with the same relative increase in theλ-derivatives. The drift and local-energy functions are easily computed to be

$$F(x) = \frac{8(1+\lambda)x^2}{5(1-\lambda)+2x^4(1+\lambda)} \tag{22}$$

$$E(x) = \frac{(1+\lambda)x^2\left(145(1-\lambda)^2 - 60(1-\lambda^2)x^4 + 4(1+\lambda)^2x^8\right)}{2(1-\lambda)\left(5(1-\lambda)+2(1+\lambda)x^4\right)^2} \tag{23}$$

In the actual simulation, these need to be evaluated at $\lambda = -2\varepsilon$, $-\varepsilon$, 0, ε, and 2ε, where $\varepsilon = 0.03\sqrt{\tau}$.

Using $10,000$ configurations, five values of τ, namely 0.10, 0.08, 0.06, 0.04, and 0.02 (it is advantageous to carry out the simulation in this order, to expedite equilibration), $40,000$ iterations at each τ (these have been divided into 5 blocks of $8,000$, yielding 5 nearly independent estimates at each τ), $a = 1.5$ in (6), $b = 0.03$ in (9) and $a = 1$ in (13), the last transformation being applied to each of the *four* companions of the leading configuration, we get results of Figure 1 (estimating the ground-state energy and its first two λ derivatives – the solid, dashed and dotted line respectively), Figure 2 (the third derivative) and Figure 3 (the fourth one). The agreement with true values is fairly good.

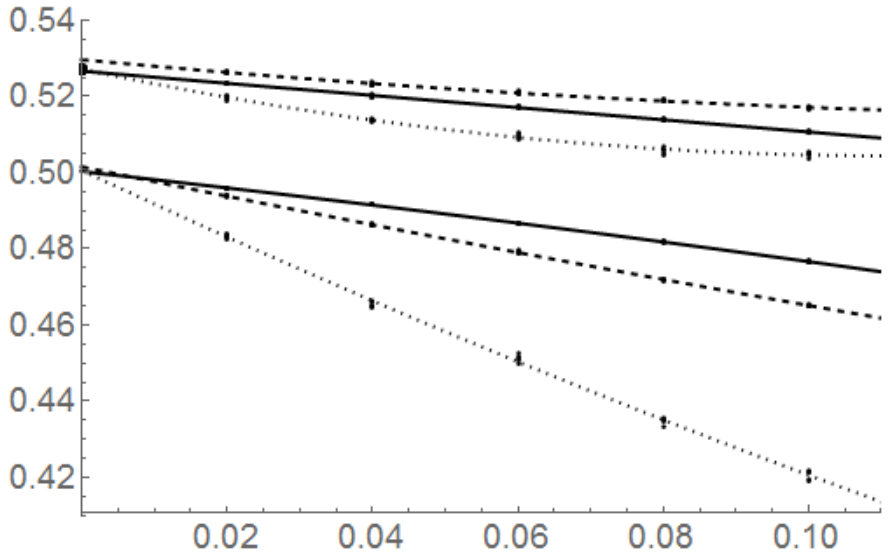

Figure 1: Variational and exact estimation of $\mathcal{E}$, $\dfrac{\mathrm{d}\mathcal{E}}{\mathrm{d}\lambda}$, and $\dfrac{\mathrm{d}^2\mathcal{E}}{\mathrm{d}\lambda^2}$.

For comparison, we also give (in the same three graphs, using dashed lines in Figure 2 and 3) the corresponding results of the variational simulation (which is achieved by simply replacing all weighed averages of the type (7) by *regular* averaging; the $S(\mathbf{x})$ quantities are then no longer needed). The target values (computed analytically) of the variational extrapolation have been indicated by half-disks attached to the vertical scale.

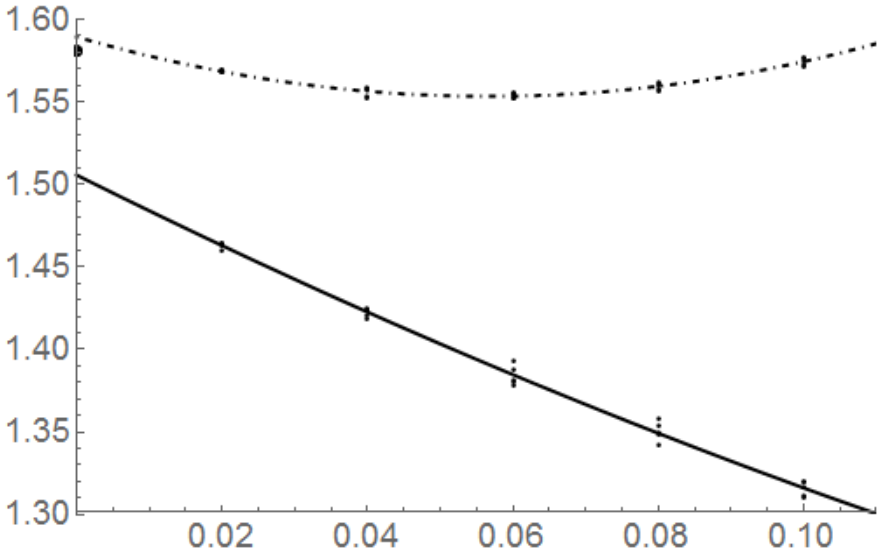

Figure 2: Variational and exact estimation of $\frac{\mathrm{d}^3\mathcal{E}}{\mathrm{d}\lambda^3}$

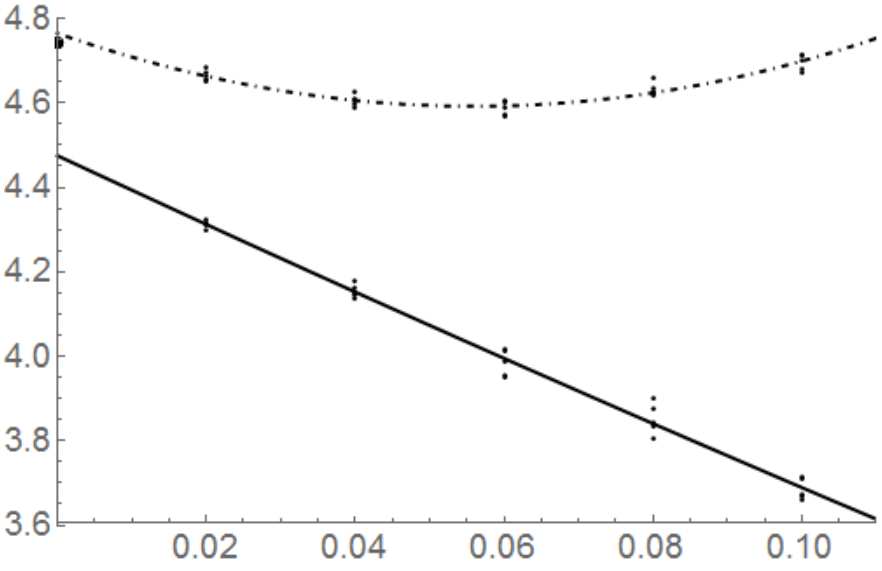

Figure 3: Variational and exact estimation of $\frac{\mathrm{d}^4\mathcal{E}}{\mathrm{d}\lambda^4}$

3 Dealing with singularities

When simulating molecules, the problem we usually encounter is that both the local energy and the drift function possess several singularities; the local energy becomes infinite (with either sign, depending on the trial function) when an electron approaches a nucleus or when two electrons closely approach each other. It is always possible to remove these singularities by a careful adjustment of trial-function parameters, but a more expedient solution is to simply truncate the values of $E(\mathbf{x})$ to stay between

$$\begin{aligned} &E_{\min} - \frac{c}{\sqrt{\tau}} \\ &E_{\max} + \frac{c}{\sqrt{\tau}} \end{aligned} \tag{24}$$

where $E_{\min}$ and $E_{\max}$ are the smallest and largest values (respectively) of the local energy function *after* excluding the *outliers* (these can be easily established from the corresponding empirical histogram), and c is suitably chosen constant. To be more specific, this means that any value of $E(\mathbf{x})$ which is bigger than the upper limit of the (24) interval is replaced by this upper limit, and similarly for the lower limit (whenever $E(\mathbf{x})$ is too small). These truncated $E(\mathbf{x})$ values are then used in the computation of $S(\mathbf{x})$ and the corresponding $W(\mathbf{x})$ weights.

There are also the (inevitable, due to Pauli exclusion principle) singularities of $\mathbf{F}(\mathbf{x})$ at the NODES (surfaces in the $\mathbf{x}$ space where $\Phi(\mathbf{x}) = 0$) of the trial function; this is a problem shared by all trial functions and can be easily dealt with by truncating all components of $\mathbf{F}$ in a manner of (24), but with its own limits.

Neither of these modifications affects the final results (each of them only introduces yet another τ-related bias which then disappears with the $\tau \to 0$ extrapolation).

A similar issue arises when computing (11) and (12); the corresponding outliers can be huge, thus ruining the resulting accuracy. To be able to deal with these, one first needs to replace the (11) and (12) estimators by

$$\overline{\frac{E_+ - E_-}{2\varepsilon}} - \tau \cdot \overline{\left(E_0 - \overline{E}_0\right) \cdot \frac{S_+ - S_-}{2\varepsilon}} \tag{25}$$

and

$$\begin{aligned} &\overline{\frac{E_+ - 2E_0 + E_-}{\varepsilon^2}} - 2\tau \cdot \overline{\left(\frac{E_+ - E_-}{2\varepsilon} - \overline{\frac{E_+ - E_-}{2\varepsilon}}\right) \cdot \frac{S_+ - S_-}{2\varepsilon}} \\ &- \tau \cdot \overline{\left(E_0 - \overline{E}_0\right) \cdot \frac{S_+ - 2S_0 + S_-}{\varepsilon^2}} + \tau^2 \cdot \overline{\left(E_0 - \overline{E}_0\right) \cdot \left(\frac{S_+ - S_-}{2\varepsilon} - \overline{\frac{S_+ - S_-}{2\varepsilon}}\right)^2} \end{aligned} \tag{26}$$

respectively, where now the bars imply averaging with respect to the W_0 weights *only*. This enables us to further truncate, individually (i.e. before any averaging), the values of

$$\frac{E_+ - E_-}{2\varepsilon},\ \frac{E_+ - 2E_0 + E_-}{\varepsilon^2},\ \tau \cdot \frac{S_+ - S_-}{2\varepsilon},\ \text{and}\ \tau \cdot \frac{S_+ - 2S_0 + S_-}{\varepsilon^2} \tag{27}$$

again in the manner of (24). This appears to solve the problem, introducing yet another τ-related and thus inconsequential bias.

The corresponding formulas for the third and higher (and also mixed, when having several parameters) derivatives can be found in [12]. To understand where these come from, just differentiate (7) with respect to λ once to get (25), twice to get (26) etc., keeping in mind that both $E(\mathbf{x}')$ and $S(\mathbf{x}')$ are functions of λ. This results in a non-trivial combination of quantities called CUMULANTS; it is the price to pay for interchanging the two operation (ensemble averaging and numerical differentiation). Utilizing cumulants in this manner appears to be yet another unique aspect of our technique.

Finally (as already mentioned), there is now the danger that one (or more) of the ensemble's triplets approaches a node so closely that one of its three companions crosses the node and the corresponding drift then sends it in *opposite* direction to the other two (as $\mathbf{F}$ always points *away* from a node); this means that the three (up to that point 'parallel') paths will suddenly start separating from each other. Even though this is now partially alleviated by treating the triplet as an outlier (since it inevitably becomes one) and truncating its contribution to our averages, we clearly cannot afford to keep on creating more of these indefinitely. The solution has already been provided by bringing the three companions of *all* triplets closer together by the (13) transformation – in addition to keeping the simulation more stable in general, this also restores 'infinitesimal' separation of any 'broken' triplet in a handful of iterations.

3.1 LiH example

In this example, we estimate the first *two* derivatives of ground-state energy with respect to λ, which is defined as the length of the distance between the two nuclei of the LiH molecule (namely Li and H, having atomic numbers of 3 and 1 respectively), initialized at $\lambda = 3$ (using atomic units). The molecule has 4 electrons, two with spin up and two with spin down; the trial function we use here has been taken from [12] and is very simple: a product of two Slater determinants built out of four atomic orbitals, further multiplied by a product of four Jastrow factors (for pairs of electrons with opposite spin). This trial function yields -8.029 as the variational estimate of the ground-state energy (whose exact value is -8.070); using (7) instead of regular averaging changes this estimate to -8.062. The reason why we are still above the correct value is due to our trial function having slightly incorrect nodes, as mentioned earlier.

For the actual simulation we used the same set of τ values and the same ε as in the harmonic-oscillator example, but this time the number of configurations was 2500, and so was the number of iterations at each τ (divided into 5 blocks of 500). Also, we had to use a new value of $a = 1/2$, $b = 0.03$, and $a' = 1/3$ (extrapolating to $\tau = 0$ becomes too difficult with higher values of a'; smaller values lead to increasing statistical error) and the following

truncations (we list the target quantity with the corresponding lower and upper limit)

Target	Lower	Upper
$E(\mathbf{x})$	$-9.2 - \dfrac{0.1}{\sqrt{\tau}}$	$-7 + \dfrac{0.1}{\sqrt{\tau}}$
$\dfrac{E_+ - E_-}{2\varepsilon}$	$-0.4 - \dfrac{0.01}{\sqrt{\tau}}$	$0.4 + \dfrac{0.01}{\sqrt{\tau}}$
$\dfrac{E_+ - 2E_0 + E_-}{\varepsilon^2}$	$-0.5 - \dfrac{0.01}{\sqrt{\tau}}$	$0.5 + \dfrac{0.01}{\sqrt{\tau}}$
$\tau \cdot \dfrac{S_+ - S_-}{2\varepsilon}$	$-1 - \dfrac{0.02}{\sqrt{\tau}}$	$1 + \dfrac{0.02}{\sqrt{\tau}}$
$\tau \cdot \dfrac{S_+ - 2S_0 + S_-}{\varepsilon^2}$	$-2 - \dfrac{0.04}{\sqrt{\tau}}$	$2 + \dfrac{0.04}{\sqrt{\tau}}$

(28)

Furthermore, each component of every $\mathbf{F}(\mathbf{x})$ was truncated to stay within the

$$\left(-3 - \frac{0.1}{\sqrt{\tau}},\, 3 + \frac{0.1}{\sqrt{\tau}}\right) \tag{29}$$

interval. All these limits have been selected based on the corresponding histogram; as an example, Figure 4 displays an ensemble's worth of $(E_+ - E_-)/(2\varepsilon)$ values (save a handful of outliers).

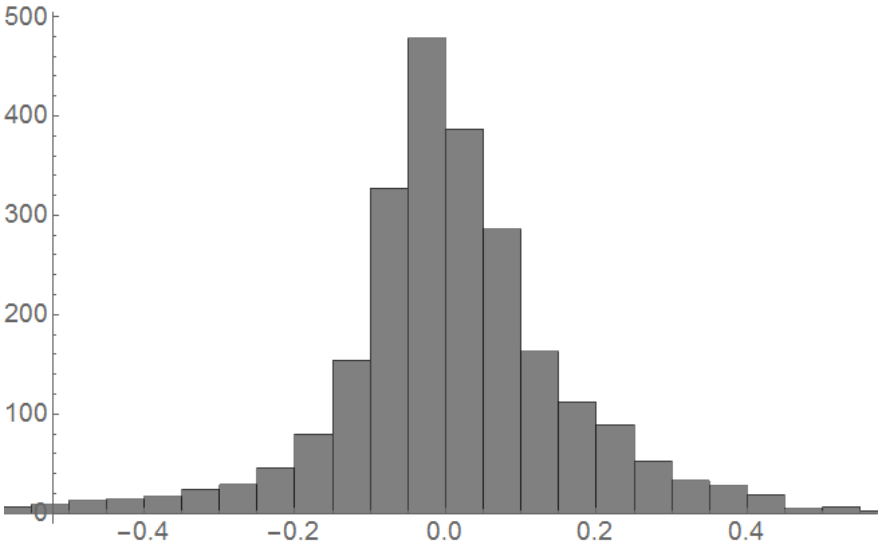

Figure 4: Histogram of $\dfrac{E_+ - E_-}{2\varepsilon}$ values.

The results are shown in Figure 6 (the first derivative, which should be close to zero) and Figure **??** (the second derivative, called the harmonic constant; its true value is about 0.0666 atomic units); numerically, our two estimates are -0.0006 ± 0.0006 and 0.0676 ± 0.0013, in good agreement with empirical values (the slightly negative first derivative implies that our $\lambda = 3$ needs to be corrected to $3 + 0.0006/0.0676 = 3.009$, in a reasonable agreement with the observed value of 3.015).

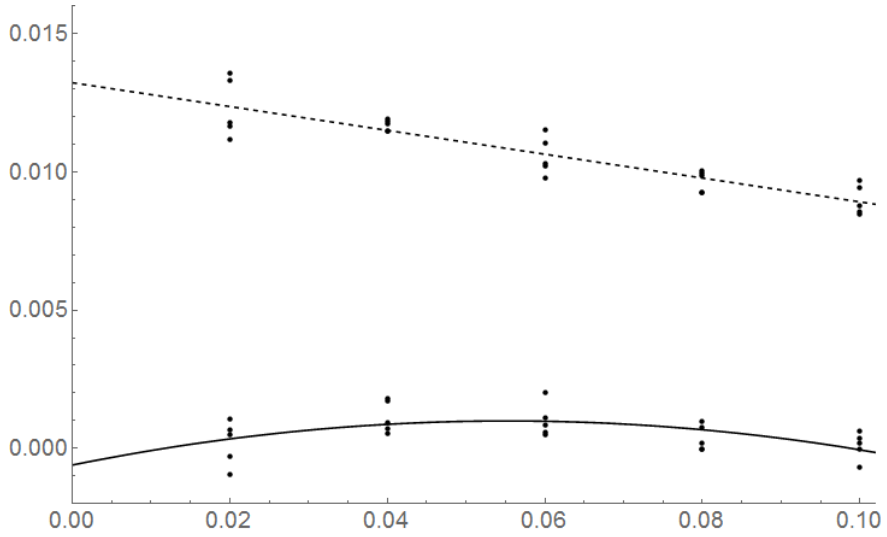

Figure 5: Variational and exact estimation of $\dfrac{\mathrm{d}E}{\mathrm{d}\lambda}$

For comparison, in the same two graphs we also show (dashed lines) the corresponding variational results, computed based on

$$\overline{\frac{E_+ - E_-}{2\varepsilon r}} \tag{30}$$

(a single-iteration estimate of the first derivative) and

$$\overline{\frac{E_+ - 2E_0 + E_-}{(\varepsilon r)^2}} \tag{31}$$

(of the second one), where the bar now indicates *regular* averaging. The resulting estimates are (not surprisingly) rather inferior.

4 Conclusion

We have demonstrated a new technique of *parallel* sampling (distinct from *correlated* sampling of existing literature) for computing derivatives of ground-state energy with respect

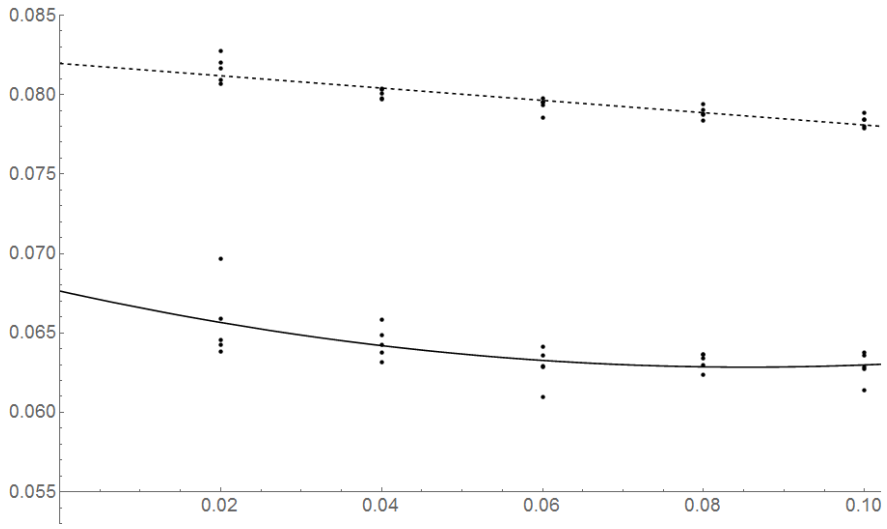

Figure 6: Variational and exact estimation of $\frac{\mathrm{d}^2 E}{\mathrm{d}\lambda^2}$

to a given parameter using a guiding function which is allowed to vary with the parameter. This is of particular importance when establishing a molecule's geometry and corresponding vibrational properties. The results are subject only to the fixed-node approximation and the inevitable statistical error. We thus claim that, in principle, one can achieve practically exact results by improving the guiding function and increasing the length of computation. This is in contrast to our previous work [12] which, to estimate the same geometry-related properties, had to introduce an extra approximation of unknown bias.

References

[1] J. B. Anderson, *J. Chem. Phys.* **63**, 1499 (1975)

[2] R. Assaraf and M. Caffarel, *J. Chem. Phys.* **119**, 10536 (2003)

[3] R. Assaraf, *Phys. Rev. E* **90**, 063317 (2014)

[4] B. M. Austin, D. Y. Zubarev and W. A. Lester Jr., *Chem. Rev.* **112**, 263 (2012)

[5] A. Badinski and R. J. Needs, *Phys. Rev. E* **76**, 036707 (2007)

[6] A. Badinski, P. D. Haynes and R. J. Needs, *Phys. Rev. B* **77**, 085111 (2008)

[7] A. Badinski and R. J. Needs, *Phys. Rev. B* **78**, 035134 (2008)

[8] C. Filippi and C. J. Umrigar, *Phys. Rev. B* **61**, R16291 (2000)

[9] W. M. C. Foulkes, L. Mitas, R. J. Needs and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001)

[10] M. H. Kalos, *Phys. Rev.* **128**, 1791 (1962)

[11] C. J. Umrigar, *Int. J. Chem. Supplement* **23**, 217 (1989)

[12] J. Vrbik and S. M. Rothstein, *J. Chem. Phys.*, **96** 2071 (1992)